

Computing News



*News from the Computing Division
Fermi National Accelerator Laboratory*

VOLUME XXIII, NUMBER 1, FEBRUARY, 1995

The Fermilab Farms in 1994

The farms in 1994 did not increase in total compute power, as no worker nodes were purchased to add to the system. The major changes occurred in the configuration of the systems. The majority of worker nodes were moved into new configurations for the CDF and D0 computing needs of RUN1B. The changes occurred rather slowly during most of 1994, because the nodes were kept in service as long as possible for fixed-target experiments while the colliders were preparing their code and the accelerator was tuning up.

The year in review

At the beginning of 1994 D0 was almost finished with their reprocessing of RUN1A data, CDF had finished with RUN1A and E665 was almost finished with the event reconstruction of RUN90 and RUN91 raw data. E706, E791 and E771 were all using the farms to reconstruct their data as quickly as possible. A small number of nodes were in CDF and D0's hands for tuning code and reconstructing small samples but the vast majority was dedicated to fixed-target use.

D0 started reconstruction of RUN1B data early in 1994 and so requested additional worker nodes early. Nodes were moved to D0 in April and again in June to bring them to 3 full "farm-

lets". Each farmlet consists of an I/O node (SGI 4D/420) with 3 ethernet interfaces and 24 worker nodes (typically SGI R3000 Indigo). The size of the D0 executable dictated that each worker node have at least 24 MB of real memory (some have 32 MB). Bottlenecks in the systems were worked on and largely solved by August. By that time the farms could keep up (barely) with the online event rate. Special projects, Monte Carlo and other requests made the balance very tenuous so a fourth set of nodes (IBM RS6000/220) was prepared and delivered to D0 by late November of 1994 (though not yet in use by the end of the year due to technical problems).

CDF began their reconstruction somewhat later in the year. There were many reasons for this — the program was not ready for production, the systems were not fully configured and functioning, and other minor problems. The I/O nodes for RUN1B were completely different from RUN1A and that change required some work and adjustments. On the SGI side of the farm an SGI Challenge XL with 4 processors, 4 ethernet interfaces and over 100 GB of disk space for output staging is being used as the I/O node, replacing an SGI 4D/420 with 2 ethernet interfaces and no staging disk. 63 worker nodes (SGI 4D/35, 16MB memory) are connected to the system. The IBM farm consists of 3 I/O nodes, RS6000/580, RS6000/530H and RS6000/590 connected to about 150 GB of disk and to IBM

What's Inside



The Fermilab Farms in 1994	1
Help Desk Questions and Answers	3
Help Desk Request Guidelines	3
Mailing Lists	4
Products	5
Unix Application Support Product News	5
Release of CERN v94b	6
Initializing PDFLIB in the CERN v94b Release	6
Unix	7
FNALU Status	7
Unix Operating System Support Status	7
Unix Shell Support	8

CDFSGA Upgrade	9
SGI Challenge L System Purchased	9
VAX	9
Bitnet and OpenVMS Alpha Systems	9
Communications	10
Data Communications and Networking WWW Access ..	10
Facility LAN Upgrade	10
Online Systems	10
Support for ACNET Information	10
CAMAC in DART	10
SAD: Simple Acquisition of (Epicure) Data	11
Computer Usage	11


RS6000/320H worker nodes with 32 MB of real memory. Additional nodes were allocated to CDF in April, July, August and October as the need to keep up with new data and to work on the backlog required more CPU power. Problems with the network configuration of the nodes caused the rampup of utilization to go much more slowly than expected. The problems were eventually tracked down and solved.

E665 finished the reconstruction of raw data from RUN90 and RUN91 in January of 1994. The allocation to E665 rapidly decreased as the only need that remained was for Monte Carlo reconstruction, which takes only a handful (8 or so) nodes. E791 finished reconstruction of raw data from RUN91 by summer of 1994. Their allocation was reduced to 2 production systems to allow them to finish any reconstruction and to handle some Monte Carlo needs. E791 is starting to use some of the IBM I/O nodes for second-pass stripping. E706 finished reconstruction of raw data but had massive Monte Carlo generation and reconstruction needs. These are being met by allocating any unused systems to them. E771 continues with approximately the same number of nodes (17) as they process their 1990 and 1991 data. E687 continues to use a small part of the farm for strips and substrips of their data.


No major conflicts in allocation have arisen during 1994. The farms were adequate for the current needs of experiments at Fermilab.

CPU utilization

The first chart below provides the summary of CPU time (in VAX-equivalent units) for the whole farm in 1994. Utilization was relatively flat in 1994, with an average of about 4400 VUPS/month. During most of the year some fraction of the farm was being reconfigured in some way to move worker nodes to CDF and/or D0. During the reconfiguration and tune-up the nodes being moved were idle and this has had an effect on the overall utilization.



The second chart shows the utilization for each of the many experiments that have used the farm. E791 was the largest user for the second year in a row. D0 offline and E706 were both quite large. CDF was about half of D0 but has been steadily increasing as the code and processing became more stable.



Plans

There seems to be no allocation crisis on the farms nor is any expected. CDF and D0 can handle the RUN1B data flow with the systems they have now (about 2/3 of the farm). The rest are being shared by the fixed-target experiments.

The current generation of farms is now beginning to age. Operating system releases are either impossible to use on the machines (due to vendor choice) or are not being put on the machine due to cost and other support factors. The SGI farms should all be at IRIX 5.2 soon and the IBM farms are at AIX 3.2.4.

The SGI 4D/25 worker nodes (the original UNIX farm) will be taken off of maintenance and allowed to fail and not be repaired (unless the repair is simple). This is not likely to have any impact. The 4D/35 worker nodes will be taken off maintenance if they can be easily fixed and/or replaced easily. Once RUN1B is finished and the final processing for the colliders is in hand much more decommissioning will be done.

The next generation of farms is now being planned. The main intent is to have computing available for the Fixed-Target run that begins in 1996. Some of the current farms will likely meet the needs of the fixed-target experiments. However, the age and cost of maintenance dictate that new farms will have to be in place sometime in 1996.

Stephen Wolbers, x3950, wolbers@fnal.gov



Help Desk Questions and Answers

Q: I am working on FNALU and do not have write access to my home directory. What is going on?

A: Most likely you have lost your Kerberos token. The file system on FNALU is AFS-based. AFS uses Kerberos authentication in order to determine if you have read and write access to directories. Normally when you login you automatically obtain a Kerberos token as part of the login process, but there are some rare circumstances when this doesn't happen. In addition, tokens have a finite lifetime (default 25 hours) so if you have been logged in for a long time, your token will expire. Use the **tokens** command to see if you have a valid token. Use the **klog** command to obtain a new token.

Q: I would like to print text files in landscape mode on a UNIX system but I can't pass VMS print queue characteristics via flpr. Is there any way that I can print in landscape?

A: The UNIX printing tools that are provided by the Computing Division require printing to a Postscript printer. There are two utilities in the **futil** product that allow you to print files in a variety of formats. The **a2ps** command takes ASCII input and produces postscript output. It provides the ability to print two forms per page, portrait and landscape, different fonts and point sizes, etc. By using UNIX pipes you can convert a file to postscript and send it to the printer. For example:

```
a2ps [a2ps options] file(s) | flpr [flpr options]
```

In addition, the **psnup** command allows you to print a postscript file; 2, 4, 8, 16, 32, or 64 forms per page. Refer to the **man** pages for **a2ps** or **psnup** for more details (% man a2ps).

Q: Why do some products work on FNALV on one day and the next I just get some error about AXP?

A: The FNALV cluster now consists of VAX nodes plus the new AXP node FNALV1. There are still some products that work on the VAXes that have not been translated to work on the new AXP node, FNALV1. While both types of computers run VMS as an operating system — and, therefore, take the same commands — they have a different internal architecture. Most products had to be rewritten to run on an AXP node. Some have not yet been converted; others never will be. If you **need** to be on a VAX node, you can choose a specific node when setting host or telnetting. If you set host or telnet to FNALV you may get FNALV1 or you might get one of the other nodes which are VAXes. To determine which node you are on enter SHOW PROCESS. If you would like to have the node name included in your prompt, include the following lines in your login.com

```
$ NODE = F$GETSYI("NODENAME")
$ SET PROMPT= "' 'NODE' $"
```

For more information about AXP, see the articles in INFO. To access INFO enter

```
$ setup info
$ info
```

To use FINGER from under AXP, enter

```
$ MFINGER username@hostname
```

Q: Why does my mail counter always say that I have more mail messages than it shows when I do a directory of new mail?

A: There's a bug in VMS which causes the mail counter to become "confused." To correct the mail counter, go into mail and enter READ/NEW as the first command within mail. This will reset the counter.

Q: What is the fastest way to get my disk quota ?

A: Send email to compdiv@fnal.gov specifying how much you have and how much you would like to have. If it is a large request, you must include justification of why you need it.

Q: I've forgotten my password (on FNAL, FNALV, FNALD, CDFSGA, FNSG01) what do I do now?

A: Prove to us that you really are who you say you are and we will reset it for you. You may either take your Fermilab ID card to Yolanda Valadez, Wilson Hall 8NE, or send email to compdiv@fnal.gov from one of your other (personal) Fermilab accounts including a phone number where you can be reached. Other methods have also been used. The key is that we must be certain that you are really the person who owns the account. Please understand, if you are having difficulty getting your password reset, that there are certain measures that must be taken to maintain the security of the Fermilab computers.

Joy Hathaway, x3649, hathaway@fnal.gov



Help Desk Request Guidelines

Times are tough, as you know. The cutbacks that you have felt have hit the Help Desk as well.

We would like to ask your cooperation in order to expedite your requests for assistance. The following information is needed in order to track down and solve your computer problems. (Yes, someday this will be a fill-in-the-blank form on the Web.) Obvious exceptions are things like print-queue resets for which we require only the queue and the system.

Information to help us reply to you.

Name:

Phone number:

Email address:

Group/Experiment:

The following information about the area where the problem occurred.

System:

Node:

Account name:

Hardware used for access:

(type of terminal, workstation, or personal computer)

User interface:

(X, Motif, Versaterm Pro) Date and time problem occurred:

Include the following information if relevant.

Software:

Device:

Log file:

Include a description of the problem.

Give a description of the problem including any possibly pertinent details such as what you were doing right before the problem occurred and whether or not it has worked for you in the past. Include a summary of any diagnostics that you have already done or that have been done for you.

Whether you're calling the help desk at x2345 or sending mail to consult (preferable for non-emergencies), please try to have this information available.

Joy Hathaway, x3649, hathaway@fnal.gov



Mailing Lists

FNAL, the lab mailserver, now supports mailing lists. A mailing list is nothing more than a collection of e-mail addresses associated by means of an alias. Many people are already users of mailing lists on Unix and VMS platforms (known as distribution lists on VMS systems).

There are many advantages to using the mail service on FNAL over lists on local systems:

- Unlike distribution lists on VMS systems, the mailing lists on FNAL will not send out headers that contain every member on the list.
- The software on the mail server provides mailing list manipulation.
- The software on the mail server provides the user with many different ways to create lists.
- The software on the mailserver can handle all different addressing platforms—smtp, bitnet, and VMSMail to name just a few.
- FNAL has been designated a critical system and is managed as such, with high uptime and availability.

List Creation

Lists can have as few names as three and as many as a few hundred. At the present time lists are created by sending a request to `postmaster@fnal.gov`. The postmaster will respond with a set of questions such as what is the name of the list, what is the list for, who will manage or moderate the list, can anyone send mail to the list or just the people on it, who is/are the moderators, and what are the e-mail addresses on the list.

The postmaster will review the list request and then create the list, test it, and notify the members of its existence.

List Manipulation

The mailing list software provides users with simple commands to gather information about lists and the ability to subscribe and unsubscribe from lists. This is accomplished by sending mail to `mailserv@fnal.gov` with the appropriate

keyword(s). The following are some of the mailing list manipulation commands:

SUBSCRIBE mail-list mail-address

This allows a user with the proper privileges to subscribe someone to a mailing list. For example, to subscribe Chuck Jones to the animation mailing list the mail message would look like:

```
To: mailserv@fnal.gov
From:
Subject:
subscribe animation Chuck_Jones@fnal.gov
```

UNSUBSCRIBE mail-list mail-address

This allows a user with the proper privileges to remove someone from a mailing list. To unsubscribe Chuck Jones to the animation mailing list the mail message would look like

```
To: mailserv@fnal.gov
From:
Subject:
unsubscribe animation Chuck_Jones@fnal.gov
```

SEND/LIST mail-list

This will send a copy of the mailing list to the requestor. To see the members of the animation mailing list:

```
To: mailserv@fnal.gov
From:
Subject:
send/list animation
```

INDEX

This command will return to the requestor a message containing all mailing list names that have been created with the INDEX privilege.

```
To: mailserv@fnal.gov
From:
Subject:
index
```

These commands will only work if the requestor has the proper permissions. If the user sends mail and gets a protection warning back and they are unsure why they should send any questions to `postmaster@fnal.gov`.

Mailing lists can prove to be very useful—from a secretary trying to send mail out about timesheets to the engineer sending notification of a new hardware update. If you have any questions or would like to have more information, please send your requests to `postmaster@fnal.gov`.

Jack Schmidt, x4060, schmidt@fnal.gov



Unix Application Support Product News

The status of products provided by Unix Application Support (UAS) is posted on a regular basis to the newsgroup `fnal.comp.unix`. In addition, this information is available via WWW at URL: <http://www-oss.fnal.gov:8000/uas/release.txt>. All Unix Application Support products are put in the UNIX KITS area which is the central distribution location for the Computing Division. The KITS area can be accessed via **upd** (specifying `fnsg01.fnal.gov` as the remote node) or via anonymous ftp to the machine `ftp.fnal.gov`.

The purpose of this article is to highlight some of the recent activities of the group.

Support for More Operating System Versions

Although the Computing Division currently supports AIX, IRIX, and SunOS/Solaris Unix-based systems, both HP-UX and OSF1 are becoming more popular and the Computing Division is evaluating these operating systems for possible full support. Thus, as appropriate, UAS is releasing HP-UX and OSF1 flavors of products. This is especially true of the core Fermi UNIX Environment (FUE).

Enhancements to UPS

UPS v3_3 is currently in beta-test. This release contains a variety of bug fixes as well as significant product enhancements. New features include:

- Extended Flavors. This feature will allow **ups** to select the product instance that best matches your operating system type and release, as well as other user specified parameters. Currently, only the operating system type is used in determining best match.
- Improved error, warning, and informational messages.
- Significantly improved error checking when modifying a **ups** database, which should eliminate the possibility of a corrupted **ups** database if only **ups** commands are used to modify the **ups** database.
- Improved handling of multiple **ups** databases when doing product declares.
- Support for listing out information about product dependencies. This feature requires **perl**. (See Support for Other Command Shells later in this article.)

Design of **ups** and **ups**-related products continues. A tentative feature list for UPS v3_4 has been developed. For more information on the status of future **ups** design refer to <http://www-oss.fnal.gov:8000/uas/ups/>.

Desktop Videoconferencing Product

As a result of several customer requests to package the desktop videoconferencing tools that will allow users to send and receive video and audio over the Internet, the UAS group has

released the product **mbone**. The **mbone** product combines into one package the binary forms of desktop videoconferencing tools such as **sd**, **vat**, **nv**, **wb**, and **imm**. The current release of the **mbone** product is v1_1. A document describing the different application tools and the networking and workstation hardware requirements is available on the OSS WWW documentation page (<http://www-oss.fnal.gov:8000/mans>) or via the Computing Division home page. The product contains manual pages for the various executables, explaining how to use each portion of the product.

Release of tcl/tk-based Products

tcl is a public domain command interpreter and **tk** is an X-toolkit on top of **tcl**. These tools have become quite popular and other packages on the Internet are being developed in these languages.

Currently, the Online Systems Department (OLS) provides some **tcl/tk**-based products in KITS for the support of **DART** and **SDSS**. UAS is preparing to take over distribution and release of these products for the general user community. In addition, the products will be provided for a larger number of UNIX based operating systems and provide a larger number of product offerings. UAS would like to express its gratitude to OLS for all the effort that has been placed in getting these products off the ground.

The specific product plans call for:

- **tcl** is the base interpreter. This will replace the current OLS product but will include both v6_7 and v7_3.
- **tk** is the X toolkit and wish interpreter. This will replace the current OLS product. UAS will provide v3_6.
- **tclX** consists of extensions to **tcl** and **wish**. This will replace the current OLS product. UAS will provide v7_3b.
- **blt** contains additions to **tk** and **wish**. This will replace the current OLS product. UAS will provide v1_7.
- **tktools** is a collection of **tk** scripts for people to use. Plans call for **tkman** (man page interface), a calendar and rolodex program, and one of the **tk** interface builders.
- **expect** provides the **exepct** and **expectk** programs which are useful in building applications that will control other programs. UAS will provide v4_8.

Availability of Other Command Shells

In July, UAS began making available two additional login shells: **tcsh** and **bash** as part of v1_0 of the **shells** product. Since that time, two other programs, **perl** and GNU **awk** (**gawk**) have been added to this product. Although these are not login shells, they are similar to the login shells as they can be used as powerful scripting languages. The latest version of the **shells** product is v1_4a.

Release of groff

groff is a freely available implementation of AT&T's Documenter Work Bench family of tools (**nroff/troff**, **tbl**, **eqn**). These tools are roughly equivalent to the **tex/latex** family and are the standard format of UNIX man pages. Not all vendors provide these tools as a part of the base operating system and thus UAS is now providing a productized version of **groff**. In

addition to the standard GNU package, a front-end for the **man** command is provided that allows **man** to use the GNU tools to format manual pages. The current release of this product is v1_09a.

Enhancements to juke

juke is a product that provides command line and C library calls to control the basic functions of a tape jukebox. Initially, the **juke** server process only ran on SGI systems. However, with release of **juke** v4_0, the **juke** server may run on AIX and SunOS 5 systems. However, you are required to install a SCSI pass-through driver (which comes with the product) on AIX and SunOS 5 systems. The **juke** package supports Exabyte and DLT tape robots.

UNIX Application Support Group, uas-group@cdcdsv0.fnal.gov



Release of CERN v94b

CERN release v94b is now available on centrally-supported platforms and in the **upd** and **VPD** KITS databases. This release was made CURRENT on Dec. 14. At that time, CERN v94a became OLD and CERN v93d was unchained in preparation for later removal.

CERN release v94b was built under the following operating systems and compilers:

Operating System	Fortran Compiler	C Compiler
AIX	xlf 2.3.000.000	xlC 1.2.001.008
IRIX 5.2	f77 4.0.1	cc 3.18
OSF1 2.0	f77 3.4	
SunOS 4.1.3	f77 2.0.1	gcc 2.5.7
SunOS 5.3	f77 3.0	cc 3.0
OpenVMS-AXP 6.1	DEC Fortran 6.2-543	DEC C 4.0-000
OpenVMS-VAX 6.1	DEC Fortran 6.2-119	DEC C 4.0-000

The IRIX kit is available in both a `-mips1` form, appropriate for machines running the R3000 processor chip, and a `-mips2` form, appropriate for machines running the R4000 chip. Also note that we have dropped support for IRIX 4 as of this release.

Due to a number of technical problems, the SunOS 4 version of CERN v94b is missing some components. Specifically, this release supports neither GKS nor Motif. Consequently, only X11 graphics will be available in v94b under SunOS.

Solaris users of **paw++** are warned that, except when working directly on the Sun workstation, it is necessary to define an environment variable named `OPENWINHOME` that gives the name of the top directory of the `openwin` directory tree. This is usually `/usr/openwin`.

As of this release of the CERN library, we terminate our support of the ULTRIX operating system. We will follow DEC and transfer our future support efforts to the OSF/1 operating system. DECstation users should retain their existing copy of the v94a release. We also terminate our support of DI3000 graphics as of this release and encourage all DI3000 users to migrate to either X11 or Motif graphics.

A known problem occurs with **paw++** under OpenVMS-VAX only: you may be logged off of your current session when you exit **paw++**. Prior to this release, VAX users were required to add a small linker option file named `VMSC.OPT` to their link list when building executables that needed several of the CERN libraries. That was to force the linker to search the C run-time library, which it did not automatically do. As of OpenVMS 6, and the introduction of DEC C, the linker does automatically search the C run-time library and, hence, adding `VMSC/OPT` to the link list is no longer necessary.

We have received a number of queries from readers of the `cern.heplib` newsgroup concerning a serious **paw** bug that was discussed at length there. The CERN development group has tracked down the origin of this bug and issued a correction for it. We have incorporated that correction in our build of release v94b, rebuilding the affected libraries and executables as necessary. Consequently, Fermilab users of this release should not experience any of the problems noted.

Lynn Garren, `x2061`, garren@fnal.gov

John Marraffino, `x4483`, marraffino@fnal.gov



Initializing PDFLIB in the CERN v94b Release

A problem has been reported with version 5 of the CERN Parton Density Function program **PDFLIB**. This is the version that accompanies release v94b of the CERN library. The manual says that the user is required to make an initializing call to the subroutine `PDFSET` and suggests that call be made with a statement like

```
CALL PDFSET('INIT0',0.0D0)
```

Two users have reported that the program crashes when this call is made this way. One of them was working on a VAX and the other on a SGI machine. Some investigation has revealed that there is a mismatch between the argument types inside and outside of the `PDFSET` routine. The workaround is to avoid calling `PDFSET` with literal arguments. Instead, make your initializing call to `PDFSET` as shown in the following code fragment.

```
CHARACTER*20  PARM(20)
DOUBLE PRECISION VAL(20)
.
.
PARM(1)='Init0'
VAL(1)=0.0D0
CALL PDFSET(PARM,VAL)
.
```

It has been verified that calling `PDFSET` in this way works properly under all operating systems we support here at Fermilab.

John Marraffino, `x4483`, marraffino@fnal.gov



Unix Operating System Support Status

FNALU Status

Several changes have been made recently to the FNALU cluster. The IBM AFS file servers have been replaced with Sun SPARCstation 5 systems. This will allow us to use disk file space more efficiently and should allow the file servers to be more easily maintainable. There are now three AFS file servers, nodes `fsus01`, `fsus03`, and `fsus04`. Node `fsus02` will be an additional file server, but has not yet been installed. The AFS fileserver changes should be completely transparent to the user. The AFS server nodes provide the filesystem for the FNALU cluster and are not part of the set of nodes that you can log in to.

Additional login nodes are being provided. A four-processor SGI Challenge XL system is being added to the cluster (node `fsgi02`) as well as a four processor Sun SPARC 20 system (node `fsui01`). Most of the standard Fermilab UNIX products will be available on the Sun node when it is released for production use. Both of these systems should be available for production by the end of January. Once the nodes are in place there will be six login nodes on FNALU, two of which are nominally "batch" machines. As utilization on the nodes grow we will develop techniques for ensuring good interactive response on the available nodes. A seventh login node, an upgraded `fnsg01`, should be added to FNALU in the next few months.

An additional 60 Gigabytes of disk storage was recently installed on the AFS file server.

Some additional software is now available on FNALU. Framemaker is now available on the SGI and Sun systems in FNALU. To use it, do a

```
setup frame
```

and then run the program by typing

```
maker
```

which will start Framemaker up and display it to whatever display your DISPLAY environment variable is set to.

We will also make available a Microsoft Windows editor and some Windows software to run under it. Windows emulation will at the least be available on the SGI nodes in the cluster by using the Softwindows emulator. Windows emulation may also be made available on the Sun system using WABI at some time in the future. Microsoft Office (Word, Excel, and PowerPoint) will be available, so you will be able to use these popular PC programs on the FNALU Unix systems. The first phase of Windows emulation on the SGI systems will be available by early February.

Finally, the ORACLE database server will be made available on FNALU some time in the next several months. Details of the ORACLE implementation are not yet set, but further information will be available soon.

Steve Hanson, x8043, hanson@fnal.gov

The Computing Division fully supports the AIX, IRIX, and Solaris (SunOS) variants of Unix. In addition, the HP-UX and OSF1 flavors of UNIX are becoming more and more popular. The purpose of this article is to update people on what is happening regarding the OS and compiler levels for each of these platforms, especially as it relates to the certification process.

AIX

The current certified versions of the Operating System is AIX 3.2.4. The certified version of the FORTRAN Compiler is 2.3 and the version of the C Compiler is 1.2.1.8. This has been the certified release level for quite some time, although IBM released AIX 3.2.5. AIX 3.2.5 is necessary for POWERPC and POWER2 architecture machines but otherwise offers no major improvements for our machines.

IBM has released AIX 4.1 which is a major OS release and will have significant differences from any of the AIX 3.2 releases. Fermilab had limited participation in the beta-program for AIX 4.1. In addition, IBM has released a new version of the FORTRAN compiler (version 3.1) and C Compiler (version 2.1). The new compilers run on both AIX 3.2 and AIX 4.1. The Computing Division has obtained an evaluation copy of the new FORTRAN compiler and ran our standard tests to verify that it supports various FORTRAN extensions that are important to us. The compiler passed all of these tests.

AIX 4.1 (which also include C version 2.1) and FORTRAN 3.1 are considered new products by IBM and thus are not provided as part of any software maintenance agreement. We are required to purchase new copies of the Operating System and compiler in order to upgrade. This requirement has greatly delayed the process of formal certification. The current farms of RS6000's will not be upgraded to AIX 4.1 because of the cost of the upgrade.

IRIX

The current certified versions of the Operating System is IRIX 5.2. The certified version of the FORTRAN Compiler is 4.0.1 and the certified version of the C compiler is 3.18. Although the majority of SGI systems at the lab have been upgraded to IRIX 5.2, a significant minority still are at IRIX 4.0.5. The Computing Division is working with its customers to upgrade to IRIX 5.2.

SGI is expected to release IRIX 5.3 in early 1995. IRIX 5.3 is the last version of the Operating System that will run on Personal Irises. IRIX 5.3 contains a lot of bug fixes, but also provides support for new hardware and a new underlying file system type. SGI has previously announced IRIX 6.0 which is a 64-bit operating system and is required for the R8000 chip which dramatically increases floating point performance. There are no R8000 based systems at the lab currently.

SunOS/Solaris

The current certified version of the Operating System is SunOS 5.3 (Solaris 2.3). The certified version of the FORTRAN and C Compiler is 3.0. Although this is the certified OS level, a majority of systems are still at SunOS 4.

Sun has just released version 3.0.1 of the FORTRAN and C Compilers which are now being evaluated. In addition, they announced SunOS 5.4 in December which is primarily incorporates bug fixes and performance enhancements. Some reports indicate a 50% improvement in performance. Evaluation of SunOS 5.4 has not yet started.

HP-UX

Recently the Computing Division has obtained several HP-UX systems in order to formally evaluate them. They are running HP-UX 9.0.5 along with a series of patches, FORTRAN compiler version 09.16 and C compiler version A.09.69. The core Fermi UNIX Environment is almost completely ported and available in KITS. Other HP-UX-based products are available in KITS as well.

OSF1

Several OSF1 systems exist within the Division and much of the FUE environment has been ported. A variety of products are available in KITS for this platform. We are currently running OSF1 2.0, FORTRAN compiler versions 3.4 and 3.5 and C compiler version 2.0. The C compiler comes bundled with the operating system and uses the same release number as the operating system.

Discussions are being held on possible upgrades to OSF1 version 3.0. Another possibility is waiting until version 3.2 becomes available (expected in early 1995).

For More Information

To follow the progress of the certification for all of these platforms, read the `fnal.comp.unix` newsgroup on a regular basis. This will contain information about what items have been tested and what problems were encountered. A review of this information should be helpful in determining what issues need to be resolved before a specific machine can and should be upgraded. Reviewing the online release notes on a system that has already been upgraded will also be quite helpful in determining the issues that will most directly impact you. In many cases, especially for major releases of the operating system, you will need to obtain new versions of any products that contain libraries and you will need to re-compile your code that uses these libraries.

After a release has been certified, contact your USS support contact for arrangements on having your system upgraded. You should also stay in contact with the support people for central systems to understand what testing facilities are available and when upgrades will be scheduled.

Matt Wicks, x8083, wicks@fnal.gov



Unix Shell Support

Background

Over the past year or so there have been several requests for the Computing Division to support additional UNIX login shells. The two shells that have been requested more than once are **tcsh**, a public domain superset of **csh** and **bash**, a GNU shell that is similar to **ksh**, with an attempt to follow the POSIX standard shell as it was defined. Recently, CDF has requested that **tcsh** be installed as the default login shell on CDFSGA and D0 has also requested that UAS officially support **tcsh**.

Quite some time ago, the infrastructure was put in place for the Fermi UNIX Environment to support both the **csh** class shells (**csh** and **tcsh**) as well as the Bourne class shells (**sh**, **ksh**, **bash**, **zsh**). Most of the necessary support for this was introduced in **ups v3_0** and **systools v4_0**. The only real obstacle for using the Bourne shell family is that some products have not been released with the appropriate setup and unsetup files for the Bourne shell class.

The most recent version of the Operating system from IBM, SGI, and Sun all provide **csh**, **ksh**, and **sh**. IRIX 5 provides **tcsh** as well. It is easy to obtain free versions of **tcsh**, **bash**, and **zsh** from the Internet. All of these shells are in use at some HEP sites.

The POSIX shell was based heavily on **ksh**. Both **ksh** and **bash** are very close to being POSIX-compliant, and thus quite similar to the POSIX shell on VMS systems. Although shell choice is somewhat a matter of personal taste, there are technical reasons why a shell such as **bash** or **ksh** is superior to **csh** and **tcsh**. On the other hand, there is a strong installed base of **csh** use and thus it is unrealistic to expect users already comfortable with **csh** to switch to a different class of shells.

Policy Statement

The Unix Application Support group (UAS) provides a UPS product, **shells**, that contains both **tcsh** and **bash**. The current script for this product copies these executables into `/usr/local/bin`. Other products such as **cmd addproduct** and **cmd shell** can easily be configured to provide support for these shells.

UAS will consider providing **zsh** as well if sufficient user demand exists. It is unlikely that UAS will provide additional shells. Requests for other shells will be considered if a compelling argument can be made.

Matt Wicks, x8083, wicks@fnal.gov



CDFSGA Upgrade

Node CDFSGA, a central CDF UNIX facility, was recently upgraded and is now an SGI Challenge XL. It now has eight 200Mhz R4400 processors, giving it total CPU power of approximately 800 MIPS. The system can be expanded to over 3000 MIPS by adding more processors. All of the peripherals from the old CDFSGA system have been migrated to the new system. The peripherals include 162 GBytes of disk, 16 dual-density 8mm tapedrives and a connection to an STK silo containing 1.2 TBytes. This migration occurred over two weekends in December, and required the cooperation of the CDF collaboration, Distributed Hardware, Data Communications, and Unix System Support. The migration went smoothly, and has resulted in CDFSGA being a much more capable system. Special thanks go to Eric Wicklund of CDF, Lisa Giacchetti of the UNIX Systems Support, and Bob Yeager of the CDF Task Force for testing products and code and making modifications quickly to get the system running again.

Some time in late January additional disk (126 GBytes) will be added to CDFSGA. This will be Fermilab's first major installation of Fast Wide SCSI on any system. Although this installation has been held up waiting for delivery of the Fast Wide 9 Gigabyte disk drives, the hardware is finally here and will be tested and installed shortly. The addition of this disk will add greatly to the storage and performance capabilities of the CDFSGA system.

Steve Hanson, x8043, hanson@fnal.gov
Steve Wolbers, x3950, wolbers@fnal.gov



SGI Challenge L System Purchased

D0 has purchased a new Silicon Graphics Challenge L system as a pilot project for performing analysis on Run Ia and Run Ib using PAW/PIAF on D0 NTUPLES. The SGI Challenge L is configured with 4 R4400 150 MHz processors, 512 Mbytes of memory, ethernet and FDDI interfaces and 170 Gbytes of SCSI-2 fast and wide disk. The Challenge L upgrade path allows for CPU, memory, high-speed I/O and SCSI-2 expansion.

The system was chosen for its high performance processing and ability to quickly access very large data sets stored locally. The D0 software port to the SGI platform also made it an attractive choice.

This will be D0's first experience with centralized, UNIX-based processing. It will initially be used as a PIAF (Parallel Interactive Analysis Facility) analysis engine, to take advantage of very high I/O bandwidth for jobs requiring access to very large data sets. Additionally, it will be evaluated for UNIX-based file serving. Other uses are sure to follow if the system performs to expectation.

The PAW (Physics Analysis Workstation) interface provides an almost invisible transition to the UNIX operating systems. Users initiate the session from their favorite local desktop (be it VAX, Alpha, Indy or PC). Minimal network traffic between the PIAF server and the PAW clients enable the analysis jobs to run equally efficiently at any remote institution as at Fermilab.

Once connected to the PIAF server via PAW, a master and four slave processes (one for each processor) are spawned, and the user initiates the transfer of their data set to the Challenge. Future plans are to have the frequently-used data sets resident and less-used data sets automatically archived on a space-needed basis. Archived data will transparently be retrieved from mass storage upon reference.

Experiences thus far with PIAF on the Challenge have demonstrated dramatic improvements over methods currently being used. As we move toward a production phase with the addition of disk and a disk management system, we anticipate exciting results.

Mary Laser, x8525, laser@fnal.gov

VMS



Bitnet and OpenVMS Alpha Systems

Bitnet is currently not available on the Alpha (AXP) nodes of the FNALV cluster. This means that usage of the SEND and RECEIVE commands (to send or receive files via Bitnet) is limited to the VAX-based nodes. However, electronic mail message can be sent to Bitnet nodes from the Alpha systems by using the FNAL Mail Server/Gateway.

Using DECnet from any VAX or Alpha OpenVMS system:

```
FNAL: :JNET% "user-id@bitnet-nodename" or
FNAL: :IN% "user-id@bitnet-nodename" preferred method
```

Using SMTP (under OpenVMS these strings would appear after SMTP% and within double quotes):

```
user-id%bitnet-nodename@FNAL.GOV or
user-id%bitnet-nodename.BITNET@FNAL.GOV
```

Further information on E-mail routing can be found in *Mail Server/Gateway User Guide* available in the Computing Library as DCD Release Note 28.

Example "To:" addresses to be used with VMS Mail to send mail to a Bitnet user (the information in uppercase should be typed as-is):

```
FNAL: :JNET% "nagy@fndcd"
FNAL: :IN% "nagy@fndcd"
SMTP% "nagy%fndcd@FNAL.GOV"
SMTP% "nagy%fndcd.BITNET@FNAL.GOV"
```

Frank Nagy, x4935, nagy@fnal.gov



Data Communications and Networking WWW Access

The Data Communications and Networking group is providing a wealth of network related information via their World Wide Web server (<http://www-dcn.fnal.gov>).

Some of the interesting areas include:

- Complete on-line maps of the network;
- Postings of scheduled network outages;
- Local- and wide-area network statistics graphs and tables. You can navigate your way through our maps and obtain statistics on network segments of interest;
- Information on MBONE (Internet desktop video conferencing) tools;
- Documents released by Data Communications and Networking group;
- Pointers to ESnet, IETF, etc. technical documents.

So stop by our server and browse through our web. We hope the information will be of interest to you and please do not hesitate to send your comments and/or suggestions to us.

Farhad Abar, x4516, abar@fnal.gov



Facility LAN Upgrade

The Data Communications group recently completed the latest upgrade to the Laboratory's LAN environment. The facility "Hub" router, a Cisco AGS+ router, has been replaced by a collapsed backbone consisting of a DEC Gigaswitch and a Cisco 7000 router. The 7000, with 5 FDDI ports, functionally replaces the AGS+ as the interconnection device between the CDF, D0, General Facility, and Computing Division LANs, providing significantly higher performance. The Gigaswitch, a 20-port FDDI switch, facilitates division of LAN FDDI backbones into multiple FDDI rings.

Prior to the upgrade, the facility's LAN architecture was one of interconnected work group LANs, each based around a single FDDI (100 Mbps) backbone. The new collapsed backbone supports multiple FDDI rings within each work group, increasing the available network bandwidth to each work group. In the cases of the CDF and D0 LANs, a single 100 Mbps FDDI work group backbone has been split into four separate 100 Mbps FDDI rings, fully interconnected at 100 Mbps. The General Facility FDDI backbone has been split into two FDDI rings, one supporting central computing resources such as FNALU, FNALV, the IBM and SGI farms, etc., while the other supports a general facility backbone extending across the site.

One other notable change is that the facility firewall router is now directly connected to the new Cisco 7000 "Hub" router,

instead of the general facility FDDI backbone. This provides more direct off-site access for the other work group LANs.

From the user's perspective, there is no discernable difference in the facility's LAN environment, with the exception of improved network performance. Since the Gigaswitch acts as an FDDI bridge, the fact that a work group "backbone" now consists of multiple FDDI rings is transparent to systems within the work group, and to the users on those systems.

The installation of an FDDI switch at the core of the facility's LAN support is the next step in an evolutionary transition toward a LAN environment based on switching technologies. Future upgrades to the facility LAN will include ethernet and FDDI switching technologies, as well as likely new LAN switching technologies such as Asynchronous Transfer Mode (ATM).

Phil DeMar, x3678, demar@fnal.gov

Online Systems



Support for ACNET Information.

As part of the DART project, software products are available to receive ACNET information and make it available to an experiments' data acquisition system. The protocol is layered on top of UDP, and receives events from the Accelerator Division front end computers. Experiments can specify the frequency of the events, and the accelerator devices listed in each type of event. The VAX/VMS version of the software receives ACNET events and enters them into a shared memory pool (called DAQ). ACNET version 1.1 is in use at D0 and at VAX-ONLINE experiments T864 and E868. The **acnet** product is available from DISTRIBUTE on BISON. The UNIX and VxWorks version of **acnet** is available from KITS, and places events in shared memory using **dfm**, DART's shared memory pool. The VxWorks version is in use at E811.

B.Joshel (AD/MA), M.Martens (AD/MI), C.Guss (E811), J.Thompson (D0), J.Streets (CD/OLS).



CAMAC in DART

The DART collaboration has produced a package of software for CAMAC access from UNIX and VxWorks platforms using several hardware interfaces. In DART, FERA and PCOS CAMAC modules are read into the fast data acquisition stream via the DYC+, however these modules need to be initialized and tested from either the host computer or a VME master. There are new products for the CES CBD8210 VME-parallel CAMAC, the Hytec VSD2992 VME-serial CAMAC and Jorway 411 SCSI-parallel and serial CAMAC branch drivers to run on the IRIX operating system. The VxWorks versions for the VME branch drivers have also been upgraded. We have

adopted the IEEE ANSI standard 758-1979 for CAMAC sub-routines, as many experiments have extant software using this interface. The VME-CAMAC drivers map the VME registers into memory and have fast (10 microsecond) CAMAC transfer time for both IRIX and VxWorks.

As part of the DART-CAMAC project, the SCSI Jorway 411 branch driver was evaluated and added to our CAMAC support. The SCSI interface is available to experiments as an upgrade to QBUS and UNIBUS Jorway 411 interfaces from PREP.

We have written a kernel level SCSI driver for the Jorway 411. The SCSI transfer overhead is of order 1 millisecond, and CAMAC data throughput is around 500 kilobytes/second. All the products now support multiple CAMAC branches and multiple users with the use of semaphores on both VxWorks and IRIX.

The CAMAC drivers are available from KITS and are called **vsd2992**, **ces8210** and **sjj**.

The DART collaboration also requires access to CAMAC from remote computers, so we have developed an RPC interface to the IEEE ANSI CAMAC routines. This has two parts, a server which runs on the host CPU which has access to the CAMAC branch driver, and a client which runs on a remote CPU. The client library contains IEEE ANSI CAMAC routine calls, and is independent of the type of CAMAC branch driver. Version 1.2 of **rcam** contains the client library for both IRIX and VxWorks, and servers for the Hytec VSD2992, CES CBD8210 and Jorway 411. The VME-CAMAC interfaces have servers for both IRIX and VxWorks.

Version 1.0 of **tcam** is a **Tcl** interface to the IEEE routines, which can be used to control either of the three branch drivers or the RPC software driver.

Finally, some experiments require fast readout of CAMAC modules which do not interface to FERA or PCOS. The Fermilab CAMAC Smart Crate Controller (SCC) can be used to read CAMAC lists at high rate (600 nanoseconds per word) and feed a DART cable with a trigger latency of around 13 microseconds. Previously, control software was only available for the SCC from VAX/VMS platforms. Version 1.0 of the **scc** product contains programs and examples to download and control the SCC from the RS232 serial port of an SGI. Support for the 68000 cross-assembler is only available from VAX/VMS.

Documentation is contained within each product, and is also available from DOCDB on WWW. Please search for CAMAC under the URL <http://fnal.a.fnal.gov:8000/docdb/> for more information about installation, programming and timings.

Dave Slimmer, x4334, slimmer@fnal.gov

Jim Meadows, x4063, meadows@fnal.gov

Carmenita Moore, x2288, carmenita@fnal.gov

Margherita Vittone, x2625, vittone@fnal.gov

Jonathan Streets, x3629, streets@fnal.gov

SAD: Simple Acquisition of (Epicure) Data

The SAD or Simple Acquisition of Data services are designed to fill the need of experimenters who wish to perform Epicure data acquisition from their experiment's computer in a simple and platform independent manner. SAD acts as an intermediary between an application on an experiment node and the data acquisition services on an Epicure node. It utilizes the RPC or Remote Procedure Call facility to implement client/server communication between a VAX/VMS system running Epicure and the experiment's platform. IRIX5 and VxWorks platforms are currently supported.

SAD comprises a simplified subset of functions for performing Epicure data acquisition. An application provides Epicure device names and reading rates, performs polling to determine when data is ready, and receives either scaled read-backs or raw data depending on the device type. All data is returned to the application once per cycle. Timestamps from the readings are optionally available. A function is also provided to translate error codes into text strings.

The SAD services require authorization for each user and client machine. For authorization, please contact Ed Dambik of EE/RD Controls at dambik@fnal.gov, or X2465. The SAD products is available through the kits distribution system, listed as **sad v1_0**. Questions concerning SAD should be forwarded to us.

Ed Dambik, x2465, dambik@fnal.gov

Dennis Black, x4525, dblack@fnal.gov

Computer Usage



The tables and the chart below summarize the computer usage of the centrally-supported systems in the months of October, November, and December 1994. The usage is displayed by platform in VUP-months.

Judith Nicholls, x3989, nicholls@fnal.gov

Computer Usage by Group, October 1994

Group	FNALV	FNALD	CDFtr	D0FS	FNALD0	CDFSGA	CLUBS	FNALU	FNSG01	Farms	Total
D0	0.3	0.0	0.0	50.0	304.6	0.0	0.7	0.2	0.5	1484.4	1840.6
E706	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	1345.4	1345.5
CDF	0.5	113.3	327.0	0.0	0.0	131.0	0.0	0.1	0.0	731.5	1303.4
E791	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	361.0	361.0
E771	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	318.7	318.8
E665	0.1	0.0	0.0	0.0	0.0	0.0	40.9	0.8	0.0	138.8	180.7
US_CMS	0.0	0.0	0.0	0.0	0.0	0.0	0.0	47.5	0.0	0.0	47.5
E683	34.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	34.9
E687	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	15.3	15.9
Acc Theory	0.0	0.0	0.0	0.0	0.0	0.0	9.4	4.0	0.0	0.0	13.4
E781	4.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	4.0
E761	0.0	0.0	0.0	0.0	0.0	0.0	3.4	0.0	0.0	0.0	3.4
BPhys	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.6	0.0	0.0	2.6
E735	1.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.2
E672	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3
E760	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1

Computer Usage by Group, November 1994

Group	FNALV	FNALD	CDFtr	D0FS	FNALD0	CDFSGA	CLUBS	FNALU	FNSG01	Farms	Total
D0	0.4	0.0	0.0	74.9	200.0	0.0	0.4	0.4	0.2	1554.8	1831.0
CDF	0.5	124.0	539.0	0.0	0.0	137.1	0.0	0.5	0.0	944.2	1745.3
E706	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	1348.2	1348.4
E771	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	327.6	327.7
E791	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	319.7	319.7
E665	0.2	0.0	0.0	0.0	0.0	0.0	40.3	1.7	0.0	151.5	193.6
US_CMS	0.0	0.0	0.0	0.0	0.0	0.0	0.0	54.3	0.0	0.0	54.3
E687	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	29.9	30.6
E683	6.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	6.6
E760	5.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	5.5
Acc Theory	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.9	0.0	0.0	3.9
E672	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1
E731	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
E789	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Computer Usage by Group, December 1994

Group	FNALV	FNALD	CDFtr	D0FS	FNALD0	CDFSGA	CLUBS	FNALU	FNSG01	Farms	Total
D0	0.3	0.0	0.0	66.5	200.0	0.0	0.1	0.3	0.1	1725.5	1992.8
CDF	0.6	128.2	365.0	0.0	0.0	193.5	0.0	0.0	0.0	985.3	1839.1
E706	0.0	0.0	0.0	0.0	0.0	0.0	43.3	0.0	0.1	1126.2	1169.6
E771	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	337.5	337.6
E791	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	274.5	274.6
E665	0.1	0.0	0.0	0.0	0.0	0.0	12.4	1.6	1.2	10.1	25.4
US_CMS	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.9	0.0	0.0	12.9
E683	8.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.6
BPhys	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.5	0.0	0.0	2.5
E687	1.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.8
E672	1.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.5
E704	1.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.1
E760	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1
E731	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
E789	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

A Guide to Computing Division Services




TELEPHONE: (708) 840-xxxx

NEW ACCOUNTS, REQUESTS, WH8X	x3205
Feynman Computing Center FAX	x3785
FAX 8th Floor	x2783
HELP DESK , helpdesk@fnal.gov	x2345
DATA CENTER SERVICES	x2746
Feynman Computing Center	x2746
DATA COMMUNICATIONS	
Trouble Reporting, svscall@fnal.gov	x3239
ONLINE SYSTEMS	
Trouble Reporting/Tech Liaison, WH8E	x3205
(On call pager through operations)	
COMPUTER HARDWARE	
Trouble Reporting	x4373
EQUIPMENT SUPPORT GROUP	
Repair Inquiry	x2688
EQUIPMENT LOGISTICS SERVICES	
Inquiry Service Counter	x3447
DIVISION MANAGEMENT	
Division Head, Joel Butler	x3148
Deputy Division Head, Vicky White	x3936
Associate Division Head, Irwin Gaines	x4022
Assistant Division Heads	
Gerry Bellendir	x3930
Judy Nicholls	x3989
Al Thomas	x3064
COMPUTER PROTECTION PROGRAM MANAGER	
Irwin Gaines	x4022
COMPUTER DIAL-UP NUMBERS AND PORT SELECTOR CLASS CODES	
Dial-In (up to 28800bps V.FC/MNP10)	840-8134
Callback (up to 28800bps V.FC/MNP10)	840-8555
FNALV VAX Cluster	FNALV
FNALD VAX Cluster	FNALD
Internet Protocol	IP
Terminal Server	LAN
DECserver, 580	DS550

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Change of Address



If you have changed your address, or are not on *Computing News* mailing list and would like to be placed on it, or would like to be removed from the list, please complete this form and mail it to the Editing Staff.

Action to be Taken (check one):

Deletion ☐

Addition ☐

Correction ☐

Return completed form to:

Pat Haring
FERMILAB - MS 120
P.O. Box 500
Batavia, IL 60510

Address Information (Include Zip Code, please)
